

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

| | | |
|---|---|---------------------------------|
| In re application of: Childress et al. | § | |
| | § | Group Art Unit: 2144 |
| Serial No. 10/753,817 | § | |
| | § | Examiner: Anwari, Maceeh |
| Filed: January 8, 2004 | § | |
| | § | |
| For: Method and Apparatus for | § | |
| Supporting Transactions | § | |

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

35525
PATENT TRADEMARK OFFICE
CUSTOMER NUMBER

APPEAL BRIEF (37 C.F.R. 41.37)

This brief is in furtherance of the Notice of Appeal, filed in this case on December 13, 2007.

A fee of \$510.00 is required for filing an Appeal Brief. Please charge this fee to IBM Corporation Deposit Account No. 09-0447. No additional fees are believed to be necessary. If, however, any additional fees are required, I authorize the Commissioner to charge these fees which may be required to IBM Corporation Deposit Account No. 09-0447. No extension of time is believed to be necessary. If, however, an extension of time is required, the extension is requested, and I authorize the Commissioner to charge any fees for this extension to IBM Corporation Deposit Account No. 09-0447.

REAL PARTY IN INTEREST

The real party in interest in this appeal is the following party: International Business Machines Corporation of Armonk, New York.

RELATED APPEALS AND INTERFERENCES

This appeal has no related proceedings or interferences.

STATUS OF CLAIMS

A. TOTAL NUMBER OF CLAIMS IN APPLICATION

The claims in the application are: 1-20

B. STATUS OF ALL THE CLAIMS IN APPLICATION

Claims canceled: none

Claims withdrawn from consideration but not canceled: 5, 6, 9, 15 and 16

Claims pending: 1-20

Claims allowed: none

Claims rejected: 1, 5-7, 9, 15-17

Claims objected to: none

C. CLAIMS ON APPEAL

The claims on appeal are: 1, 7 and 17

STATUS OF AMENDMENTS

No amendment after final rejection was filed for this case.

SUMMARY OF CLAIMED SUBJECT MATTER

Businesses and other organizations employ network data processing systems to conduct business and other transactions. These networks may be as small as a single LAN or may encompass many networks, including the Internet. Enterprise networking involves using a network infrastructure in a large enterprise or business organization with multiple computer systems and networks. These types of infrastructures are typically extraordinarily complex. An enormous amount of effort goes into planning and managing the integration of different disparate networks and systems. Also, planning for additional interfaces as needs and demands change also occurs.

In managing an enterprise system, these systems often include a number of servers that are assigned to provide different services. Management of these servers is an important function of ensuring that services are provided when needed. Managing the allocation of resources for providing services to process requests is an important and complex task. As part of a process to identify the capability and usage of resources, identifying transactions processed by nodes, such as servers, is important for use in ensuring that a perceived capability matches the actual usage for those nodes. For example, a set of servers may be provisioned to handle requests for a Website set up to support an online business that provides goods or services. The servers also may be set up to provide access to data, such as medical records, tax information, or regulations. The resources needed vary depending on the usage and demand from clients. In provisioning resources, it is important to identify the usage of the resources. If the usage increases, capacity may be added to meet the increasing demand. In some cases, the addition of servers may be unnecessary because one or more current servers may be underutilized while others may be strained to the point of failure or are unable to meet expected service levels. A mismatch in the capabilities is often identified by the occurrence of a failure and subsequent analysis of the system. These failures typically occur when currently used load balancing techniques are unable to adequately monitor and maintain the capabilities for servicing requests.

When an application is simple and does not require the state to persist over multiple requests from a user, the normal round robin or other such load balancing techniques are sufficient to maintain capabilities for servicing requests. In the case where the application is more complex and requires state information to persist across multiple requests, the presently

available load balancing techniques are unable to sufficiently monitor and manage resources for servicing requests. In the case where state information is persisted, the user's session is required to be associated with a particular server providing the information. This situation is generally referred to as "sticky load balancing". In this case it is normal for a single server to become overloaded due to the stickiness of the transaction. This problem increases when the situation changes from the user being a human using a browser to a computer using Web services. The main reason for having to maintain state information in these examples is the need to access legacy systems.

Generally provided by the present pending claims is a technique for monitoring and identifying transactions handled by nodes in a networked data processing system, in order to provide selective load-balancing which is non-invasive in that the node does not have to respond to direct requests for the transaction information. This non-invasive monitoring is particularly useful when a given server has turned off or otherwise disabled any ability to monitor the server, such as when the server is operating during a critical time period or is otherwise heavily loaded, and therefore cannot spare the additional computing resources/overhead that would otherwise be required to provide monitoring/status information by such server.

A. CLAIM 1 - INDEPENDENT

The subject matter of Claim 1 is directed to a method in a data processing system for monitoring transactions for a set of known nodes in a network data processing system (Specification page 12, lines 15-17). Cache data from a router in the data processing system is received, where the cache data includes an identification of the set of known nodes sending data packets for transactions onto the network data processing system (Specification page 17, lines 25-26; page 14, lines 14-24; Figure 8, block 800; Figure 5, all blocks). The transactions handled by each node in the set of known nodes are identified using the cache data received from the router, to form identified transactions (Specification page 17, lines 28-29; Figure 8, block 804). The identified transactions are analyzed, and in response to analyzing the identified transactions, a load balancing process is selectively initiated for at least one of the nodes in the set of known nodes to mitigate transaction overload at the at least one of the nodes (Specification page 18, line 4 – page 19, line 26; Figure 8, block 808; Figure 9, all blocks).

B. CLAIM 10 - INDEPENDENT

The subject matter of Claim 10 is directed to a data processing system for monitoring transactions for a set of known nodes in a network data processing system (Specification page 12, lines 15-17), the data processing system comprises a bus system, a communications unit connected to the bus system, a memory connected to the bus system, and a processing unit connected to the bus system (Specification page 10, line 26 – page 12, line 14; element 300 of Figure 3). The processing unit executes a set of instructions in the memory (i) to receive cache data from a router in the data processing system, where the cache data includes an identification of the set of known nodes sending data packets for transactions onto the network data processing system (Specification page 17, lines 25-26; page 14, lines 14-24; Figure 8, block 800; Figure 5, all blocks), (ii) identifies the transactions handled by each node in the set of known nodes using the cache data received from the router (Specification page 17, lines 28-29; Figure 8, block 804); (iii) analyzes the identified transactions (Specification page 18, line 4 – page 19, line 3; Figure 8, block 808); and (iv) in response to the analyzing the identified transactions, selectively initiates a load balancing process for at least one of the nodes in the set of known nodes to mitigate transaction overload at the at least one of the nodes (Specification page 19, lines 4-26; Figure 9, all blocks).

C. CLAIM 11 - INDEPENDENT

The subject matter of Claim 11 is directed to a data processing system, including a data processor, for monitoring transactions for a set of known nodes in a network data processing system (Specification page 12, lines 15-17). The data processing system comprises receiving means for receiving cache data from a router in the data processing system, where the cache data includes an identification of the set of known nodes sending data packets for transactions onto the network data processing system (Specification page 17, lines 25-26; page 14, lines 14-24; Figure 8, block 800; Figure 5, all blocks); identifying means for identifying the transactions handled by each node in the set of known nodes using the cache data received from the router, to form identified transactions (Specification page 17, lines 28-29; Figure 8, block 804); analyzing means for analyzing the identified transactions (Specification page 18, line 4 – page 19, line 3; Figure 8, block 808); and initiating means for selectively initiating, responsive to the analyzing

means for analyzing the identified transactions, a load balancing process for at least one of the nodes in the set of known nodes to mitigate transaction overload at the at least one of the nodes (Specification page 19, lines 4-26; Figure 9, all blocks).

The equivalent structure for each of the receiving means, identifying means, analyzing means and initiating means is provided by element 300 of Figure 3, as described in the Specification at page 10, line 26 – page 12, line 14.

D. CLAIM 18 - INDEPENDENT

The subject matter of Claim 18 is directed to a computer readable medium encoded with a computer program product that is operable in a data processing system for monitoring transactions for a set of known nodes in a network data processing system (Specification page 12, lines 15-17; page 20, lines 13-28). The computer program product comprises first instructions for receiving cache data from a router in the data processing system, wherein the cache data includes an identification of the set of known nodes sending data packets for transactions onto the network data processing system (Specification page 17, lines 25-26; page 14, lines 14-24; Figure 8, block 800; Figure 5, all blocks); second instructions for identifying the transactions handled by each node in the set of known nodes using the cache data received from the router, to form identified transactions (Specification page 17, lines 28-29; Figure 8, block 804); third instructions for analyzing the identified transactions (Specification page 18, line 4 – page 19, line 3; Figure 8, block 808); and fourth instructions for selectively initiating, in response to the third instructions for analyzing the identified transactions, a load balancing process for at least one of the nodes in the set of known nodes to mitigate transaction overload at the at least one of the nodes (Specification page 19, lines 4-26; Figure 9, all blocks).

GROUND OF REJECTION TO BE REVIEWED ON APPEAL

The grounds of rejection to review on appeal are as follows:

A. GROUND OF REJECTION 1

Whether the Examiner failed to state a *prima facie* obviousness rejection under 35 U.S.C. § 103 against Claim 1 over Nelson et al. (US Publication No. 2003/0005092 A1, hereinafter ‘Nelson’) in view of Datta et al. (U.S. Patent No. 6,295,276, hereinafter ‘Datta’).

B. GROUND OF REJECTION 2

Whether the Examiner failed to state a *prima facie* obviousness rejection under 35 U.S.C. § 103 against Claims 7 and 17 over Nelson et al. (US Publication No. 2003/0005092 A1, hereinafter ‘Nelson’) in view of Nelson et al. (Patent No. 5,835,720, hereinafter ‘Nelson II’).

ARGUMENT

A. GROUND OF REJECTION 1 (Claim 1)

Claim 1 stands rejected under 35 U.S.C. § 103 as being obvious over Nelson et al. (US Publication No. 2003/0005092 A1, hereinafter, Nelson) in view of Datta et al. (U.S. Patent No. 6,295,276, hereinafter ‘Datta’).

1. Claim 1

With respect to Claim 1, it is urged that the cited Nelson reference does not describe any type of *identification with respect to transactions*, but instead is directed to a technique for *identifying physical hardware devices* (Nelson paragraphs [0002] and [0006]). Quite simply, a transaction (as claimed) is not equivalent to a device¹.

In addition, Nelson does not provide any type of *load balancing as a result of any transaction identification*, but instead provides a user with a *list of discovered devices* (Nelson paragraphs [0011] and [0051]; Figure 2, block 216). The Nelson system is not directed to performance monitoring via transaction identification and subsequent analysis thereof, and resulting load-balancing that is done to mitigate transaction overload conditions, but instead is directed to a technique for identifying physical devices on a network to assist law enforcement personnel in recovery efforts associated with stolen property (Nelson paragraph [0002]).

In addition, and contrary to the Examiner’s assertion, Nelson’s “range walk” discovery technique described in paragraphs [0006], [0011] and [0026] does not teach or otherwise describe a *node load balancing technique*, but instead describes a technique for locating and identifying devices.

Nor is the Nelson “range walk” done *in response to analyzing identified transactions*, as the Nelson “range walk” is what is used as a *precursor step prior to any type of device identification* (Nelson block 204 of Figure 2, which performs the “range walk”, is done *before any device*

¹ transaction: the act or process of transacting, The American Heritage Desk Dictionary, Houghton Mifflin Company, copyright 1981.

device: something devised for a particular purpose, esp., a machine, The American Heritage Desk Dictionary, Houghton Mifflin Company, copyright 1981.

identification that occurs in block 206 of Figure 2 since the devices have not yet been discovered/located – which is what the “range walk discovery technique” is all about). Restated, Nelson’s “range walk” is done to identify devices, and is not done in response to identifying anything (either identified transactions, as claimed, or any other type identification). In contrast, per the features of Claim 1, the *selectively initiating of the load balancing* process is done *in response to analyzing the identified transactions*.

Nor do the teachings of the cited Datta reference overcome such deficiency. The Examiner cites Datta col. 15, lines 16-24 and Figures 4-5 in rejecting Claim 1. Appellants urge that this cited passage merely describes that traditional load balancing techniques are used by a router selector to *select a router using load information 410*. Importantly, this router selection using traditional load balancing techniques is not described as being done *in response to analyzing the identified transactions*, as required by the features of Claim 1. Instead, the Datta load balancing occurs in response to a client device’s *request to access a resource* (Datta col. 4, lines 26-41). Thus, it is further urged that a proper prima facie case of obviousness has not been established as none of the cited references teach or suggest (i) *identifying transactions handled by each node* in a set of nodes, (ii) *analyzing these identified transactions* that are handled by each node, or (iii) *selectively initiating a load balancing technique in response to the analyzing of the identified transactions* handled by each node – which advantageously allows for providing selective load-balancing which is non-invasive in that the node does not have to respond to direct requests for the transaction information and yet nodal transaction information is still able to be used to facilitate load balancing.

Still further, the Examiner is using impermissible hindsight analysis in making the Nelson/Datta combination. The Examiner states that it would have been obvious to add Datta’s load balancing to Nelson ‘to avoid overloading any individual remote node’. Appellants urge clear error, as Nelson does not describe any individual remote nodes that are susceptible to overloading and therefore there would be no reason to add a node load balancing technique to Nelson as the problems for which such load balancing may solve do not in fact exist in the simplified Nelson device discovery system that *does not share or balance system resources*. Thus, the only reason for adding nodal load balancing to the teachings of Nelson must instead be coming from Appellants’ own disclosure and claims, which is impermissible hindsight analysis².

² It is error to reconstruct the patentee’s claimed invention from the prior art by using the patentee’s

Thus, the cited references have been improperly combined using impermissible hindsight analysis. Even when improperly combined, there are still missing claimed features that are not taught or suggested by the resulting combination – strongly evidencing non-obviousness. Accordingly, it is urged that the Examiner has therefore failed to properly establish a prima facie showing of obviousness with respect to Claim 1, and therefore Claim 1 has been erroneously rejected under 35 U.S.C. § 103³.

B. GROUND OF REJECTION 2 (Claims 7 and 17)

Claims 7 and 17 stand rejected under 35 U.S.C. § 103 as being unpatentable over Nelson et al. (US Publication No.: 2003/0005092 A1) in view of Nelson et al. (US Patent No.: 5,835,720).

1. Claims 7 and 17

The rejection of Claim 7 (and similarly for Claim 17) is respectfully traversed for similar reasons to those given above with respect to Claim 1, as the newly cited US Patent 5,835,720 to Nelson does not overcome the teaching/suggestion deficiencies, and the associated failure of the Examiner to properly establish a prima facie showing of obviousness, identified above with respect to Claim 1.

Further yet, Claim 7 depends upon Claim 1 and thus includes all of the features recited in Claim 1. In rejecting Claim 1, the Examiner relies upon the cited Datta reference, and yet in the rejection of Claim 7 (which depends upon Claim 1) the Datta reference is not relied upon – thus, whatever Datta teachings were relied upon in rejecting Claim 1 are clearly not taught by the Nelson/Nelson II teachings, further evidencing an improper prima facie showing of obviousness with respect to Claim 7 (and similarly for Claim 17) – and in particular the cited Nelson/Nelson

claims as a “blueprint”. When prior art references require selective combination to render obvious a subsequent invention, there must be some reason for the combination other than the hindsight obtained from the invention itself. *Interconnect Planning Corp. v. Feil*, 774 F.2d 1132, 227 USPQ 543 (Fed. Cir. 1985).

³ To establish prima facie obviousness of a claimed invention, *all of the claim limitations must be taught or suggested by the prior art*. MPEP 2143.03 (emphasis added by Appellants). *See also, In re Royka*, 490 F.2d 580 (C.C.P.A. 1974). If the examiner fails to establish a prima facie case, the rejection is improper and will be overturned. *In re Fine*, 837 F.2d 1071, 1074, 5 USPQ2d 1596, 1598 (Fed. Cir. 1988).

II combination does not teach or otherwise suggest (nor has the Examiner even alleged a teaching/suggestion of) “in response to the analyzing the identified transactions, selectively initiating a load balancing process for at least one of the nodes in the set of known nodes to mitigate transaction overload at the at least one of the nodes”.

Thus, it is further shown that Claim 7 (and similarly for Claim 17) has been erroneously rejected due to this additional prima facie obviousness deficiency.

C. CONCLUSION

As shown above, the Examiner has failed to state valid rejections against any of the claims. Therefore, Appellants request that the Board of Patent Appeals and Interferences reverse the rejections. Additionally, Appellants request that the Board direct the examiner to allow the claims.

/Wayne P. Bailey/

Wayne P. Bailey

Reg. No. 34,289

YEE & ASSOCIATES, P.C.

PO Box 802333

Dallas, TX 75380

(972) 385-8777

CLAIMS APPENDIX

The text of the claims involved in the appeal is as follows:

1. A method in a data processing system for monitoring transactions for a set of known nodes in a network data processing system, the method comprising:

receiving cache data from a router in the data processing system, wherein the cache data includes an identification of the set of known nodes sending data packets for transactions onto the network data processing system;

identifying the transactions handled by each node in the set of known nodes using the cache data received from the router, to form identified transactions;

analyzing the identified transactions; and

in response to the analyzing the identified transactions, selectively initiating a load balancing process for at least one of the nodes in the set of known nodes to mitigate transaction overload at the at least one of the nodes.
2. The method of claim 1, wherein the cache data is from an address resolution protocol cache located on the router.
3. The method of claim 1 further comprising receiving cache data from other routers on the network data processing system.
4. The method of claim 1, wherein the receiving step occurs on a periodic basis.

7. The method of claim 4 further comprising:
- generating a display of the set of known nodes in a graphical view, wherein the graphical view includes the communications paths with a graphical indication of the network traffic.
8. The method of claim 2, wherein the cache data is received through an agent located on the router.
10. A data processing system for monitoring transactions for a set of known nodes in a network data processing system, the data processing system comprising:
- a bus system;
 - a communications unit connected to the bus system;
 - a memory connected to the bus system, wherein the memory includes a set of instructions; and
 - a processing unit connected to the bus system, in which the processing unit executes the set of instructions to receive cache data from a router in the data processing system, in which the cache data includes an identification of the set of known nodes sending data packets for transactions onto the network data processing system, identifies the transactions handled by each node in the set of known nodes using the cache data received from the router, to form identified transactions; analyzes the identified transactions; and in response to the analyzing the identified transactions, selectively initiates a load balancing process for at least one of the nodes in the set of known nodes to mitigate transaction overload at the at least one of the nodes.
11. A data processing system, including a data processor, for monitoring transactions for a set of known nodes in a network data processing system, the data processing system comprising:
- receiving means for receiving cache data from a router in the data processing system, wherein the cache data includes an identification of the set of known nodes sending data packets for transactions onto the network data processing system;

identifying means for identifying the transactions handled by each node in the set of known nodes using the cache data received from the router, to form identified transactions;

analyzing means for analyzing the identified transactions; and

initiating means for selectively initiating, responsive to the analyzing means for analyzing the identified transactions, a load balancing process for at least one of the nodes in the set of known nodes to mitigate transaction overload at the at least one of the nodes.

12. The data processing system of claim 11, wherein the cache data is from an address resolution protocol cache located on the router.

13. The data processing system of claim 11 wherein the receiving means is a first receiving means and further comprising:

second receiving means for receiving cache data from other routers on the network data processing system.

14. The data processing system of claim 11, wherein the receiving means is initiated on a periodic basis.

17. The data processing system of claim 14 further comprising:

generating means for generating a display of the set of known nodes in a graphical view, wherein the graphical view includes the communications paths with a graphical indication of the network traffic.

18. A computer readable medium encoded with a computer program product that is operable in a data processing system for monitoring transactions for a set of known nodes in a network data processing system, the computer program product comprising:

first instructions for receiving cache data from a router in the data processing system, wherein the cache data includes an identification of the set of known nodes sending data packets for transactions onto the network data processing system;

second instructions for identifying the transactions handled by each node in the set of known nodes using the cache data received from the router, to form identified transactions;

third instructions for analyzing the identified transactions; and

fourth instructions for selectively initiating, in response to the third instructions for analyzing the identified transactions, a load balancing process for at least one of the nodes in the set of known nodes to mitigate transaction overload at the at least one of the nodes.

19. The computer program product of claim 18, wherein the cache data is from an address resolution protocol cache located on the router.

20. The computer program product of claim 18 further comprising:

fifth instructions for receiving cache data from other routers on the network data processing system.

EVIDENCE APPENDIX

This appeal brief presents no additional evidence.

RELATED PROCEEDINGS APPENDIX

This appeal has no related proceedings.